

REMARKS

Reconsideration of the present application is respectfully requested. In this amendment, claims 9, 18 and 21 have been amended, and claims 20, 22 and 23 have been canceled. Claim 17 was previously canceled. No new matter has been added.

Summary of Rejections

Claims 1, 2 7-11, 16, 18, 21 and 22 stand rejected under 35 U.S.C. § 103(a) based on U.S. Patent no. 6,505,205 of Yanai et al. ("Yanai") in view of U.S. Patent no. 6,728,898 of Tremblay et al. ("Tremblay"). Claims 3-6, 12-15, 20 and 23 stand rejected under 35 U.S.C. § 103(a) based on Yanai in view of Tremblay, and further in view of U.S. Patent Application Publication no. 2002/0178355 of Selkirk et al. ("Selkirk").

Response to Rejections

Before discussing the cited art and the claims, an overview and some background regarding the present invention may be helpful. The invention generally relates to a technique for mirroring data using two separate storage systems (e.g., storage servers) on a network, i.e., a source storage system and a destination storage system. The source storage system receives write requests from clients (for modifying stored data); however, the *destination* storage system does *not* receive the write requests from clients (see, e.g., Fig. 1 of the present application). The destination storage system determines how its mirror copy of data should be updated based only on information it receives from a source storage system. In other words, the destination receives all modified data only from a source storage system. The modified data can be

sent from the source storage system to the destination storage system in the form of log entries that reflect individual write transactions which the source storage system has received from clients (see, e.g., Fig. 6, blocks 605 – 607, of the present application).

Periodically, or according to some schedule, any data recently modified at the source storage system is saved to disk (or other persistent storage); this event is called a "consistency point" (CP) (see, e.g., present application at paragraph [0005]). Part of a CP event is updating the data at the destination storage system to reflect any recent changes, which is referred to as the "synchronization" phase of the CP event, or "CP sync".

In the known prior art, the CP sync phase involved sending over the network, from the source system to the destination system, all of the data modified since the last CP. In a system which handles large volumes of data, that would consume significant amounts of network bandwidth. See present application at paragraph [0006].

The present invention overcomes this problem. In particular, at least some of the modified data is not required to be sent from the source to the destination during the CP sync phase. Instead, only references to such data are sent to the destination during the CP sync phase. See, e.g., present application at paragraph [0020] et seq.

The present application contains six (6) independent claims, namely, claims 1, 9, 13, 16, 18 and 21. Claim 1, for example, recites:

1. (Currently amended) A method of mirroring data stored in a source storage system, the method comprising:
receiving at the source storage system a plurality of requests from a set of clients, the requests indicating modifications to be made to stored data;

saving modified data in the source storage system based on the requests;

receiving the modified data at a destination storage system from the storage system, **wherein the destination storage system is configured to receive the modified data from the source storage system and not from any client of the set of clients;** and

during a synchronization phase, synchronizing data stored in the destination storage system with the data stored in the source storage system, including mirroring at least a portion of the modified data in the destination storage system **without requiring said portion of the modified data to be sent from the source storage system to the destination storage system during the synchronization phase.** (Emphasis added.)

The Examiner admits that Yanai does not teach “mirroring at least a portion of the modified data in the destination storage system without requiring said portion of the modified data to be sent from the source storage system to the destination storage system during the synchronization phase,” per claim 1. (Final Office Action, p. 3.)

However, the Examiner contends that Tremblay discloses this feature and contends that it would be obvious to combine this (alleged) teaching of Tremblay with the teachings of Yanai.

Applicant respectfully submits that no combination of the cited references discloses the method of claim 1 or renders it obvious. Applicant respectfully submits that the proposed combination of Yanai and Tremblay is improper, as discussed below. Yanai expressly *teaches away* from such a combination with Tremblay. Further, the alleged *motivation* to combine Yanai and Tremblay does not withstand scrutiny.

A. Yanai teaches away from combination with Tremblay.

Yanai expressly *teaches away* from such a combination with Tremblay.

Tremblay discloses a method of mirroring data within a particular computer system, where *the source and destination storage system both receive write data from the host* (col. 3, lines 42-43). Yanai, on the other hand, expressly points out the *disadvantages* of such an approach:

Some providers of prior art data storage systems have proposed a method of data mirroring *whereby one host Central Processing Unit (CPU) or processor writes data to both a primary, as well as a secondary, data storage device or system*. Such a proposed method, however, *overly burdens the host CPU with the task of writing the data to a secondary storage system and thus dramatically impacts and reduces system performance*. Yanai, col. 2, lines 14-21 (emphasis added).

It is apparent that the above-quoted section from Yanai describes at a high level the basic approach disclosed in Tremblay (i.e., where source and destination both receive write requests from the host) and points out its shortcomings. In the "Summary of the Invention" and "Detailed Description," Yanai then goes on to describe embodiments in which the host does *not* write to both the source and destination storage system, as an improvement over the above-noted technique. Therefore, Yanai clearly *teaches away* from the proposed combination with Tremblay.

B. Alleged *motivation* to combine does not withstand scrutiny.

The Examiner contends that the motivation to combine the teachings of Tremblay with those of Yanai would be "confirming data associated with all write requests that preceded the designated write request have been written to the secondary storage (column 1, lines 46-67)." (Final Office Action, page 4.) However, Yanai *already has* the

capability of “confirming data associated with all write requests that preceded the designated write request have been written to the secondary storage.” This is evident in Yanai’s description of the “synchronous mode”, in which the primary storage system does not acknowledge completion of a write to the host until the data has been successfully written to the secondary storage system (see col. 3, lines 32-41, and col. 15, line 65 – col. 16, line 4). Since Yanai already has that capability, there would be no reason for one to try to combine the cited teaching in Tremblay with Yanai. For at least this reason, the motivation alleged by the Examiner does not withstand scrutiny.

Second, the alleged motivation also does not make sense, because contrary to the Examiner’s position, the feature/characteristic being cited in Tremblay (i.e., mirroring at least a portion of the modified data in the destination storage system without requiring said portion of the modified data to be sent from the source storage system to the destination storage system during the synchronization phase) is *not for the purpose of* “confirming data associated with all write requests that preceded the designated write request have been written to the secondary storage.” Rather, in Tremblay, the fact that modified data is mirrored in the destination storage system “without requiring it to be sent from the source storage system during synchronization” is just an inherent aspect of the system architecture. Because the client sends the same write requests to both I/O controllers in Tremblay, (col. 3, lines 42-43), *there is no need to send any of the write data from source to destination when a synchronization is needed*. The cited feature/characteristic in Tremblay has nothing to do with confirming that write requests have been written to storage, and as such, it is not done for the purpose which the Examiner contends. For this additional reason, therefore, the alleged motivation to

combine the teachings of Tremblay with those of Yanai does not withstand scrutiny. Indeed, it is clear that there would be *no motivation* to combine the cited teachings of Tremblay and Yanai.

Thus, because Yanai expressly *teaches away* from the proposed combination with Tremblay, and because there would be no *motivation* to combine their teachings, the proposed combination of Yanai and Tremblay is improper. And, because all of the rejections depend on that combination, all of Applicants' claims are patentable over the cited art.

Independent claims 9, 13, 18 and 21

Applicants' arguments above apply to *all* of Applicants' independent claims. In addition, further regarding independent claims 9, 13, 18 and 21, Applicants respectfully submit that Examiner's interpretation of the cited art is inaccurate.

Note that claims 9, 13, 18 and 21 incorporate limitations similar to limitations in dependent claims 2-4. Claim 9, for example, now recites the following limitations:

[F]or each block of a second subset of the plurality of blocks, sending a reference from the first storage appliance to the second storage appliance, instead of sending the corresponding block, **each said reference identifying a data transfer during which the corresponding block was previously sent from the first storage appliance to the second storage appliance and a location of said block within the data transfer, each said reference for use by the second storage appliance to locate the corresponding block in local storage of the second storage appliance and to store the corresponding block in the second set of non-volatile storage devices.** (Emphasis added.)

Note also that claims 1 and 2 were rejected based on a combination of Yanai/Tremblay, while claims 3 and 4 were rejected based on a combination of Yanai/Tremblay/Selkirk.

As a preliminary matter, Applicants would like to point out that the Examiner did not provide any rationale for rejecting dependent claim 2. The Examiner merely lumped claim 2 into the same overall basis of rejection (Yanai/Tremblay) as claim 1, without further comment. The Examiner is reminded that "[i]t is important for an examiner to properly communicate the basis for a rejection so that the issues can be identified early and the applicant can be given a fair opportunity to reply." MPEP § 706.02(j). By not specifically explaining the rejection of claim 2, the Examiner has failed to establish a *prima facie* case of obviousness regarding that claim.

Further, one of the requirements to establish a *prima facie* case of obviousness is that the cited references must teach or suggest *all of the claim limitations*. *In re Vaeck*, 947 F.2d 488, 20 USPQ.2d 1438 (Fed. Cir. 1991); MPEP § 706.02(j). Applicants find no disclosure or suggestion, in any of the cited references, of mirroring at least a portion of the modified data in the destination storage system without requiring said portion of the modified data to be sent from the source storage system to the destination storage system during the synchronization phase, by using references to said portion of modified data, said references being received at the destination storage system from the source storage system during the synchronization phase, *each said reference identifying a data transfer during which a corresponding block was previously sent from the source storage system to the destination storage system and a location of said block within the data transfer, wherein the destination storage system uses said*

references to locate corresponding blocks of modified data in storage associated with the destination storage system.

As noted above, some of these limitations are similar to the limitations in dependent claims 3 and 4. In rejecting claim 3, the Examiner contends that “Selkirk teaches a method, wherein each reference comprises a transfer ID indicating a data transfer in which the corresponding block was previously sent from the source storage system to the destination storage system,” citing Selkirk at page 5, paragraph 85, and page 6, paragraphs 86 and 87 (Final Office Action, p. 6). Applicants respectfully disagree. Applicants find no disclosure, in the cited section or elsewhere in Selkirk, of a transfer ID indicating a data transfer in which the corresponding block was previously sent from the source storage system to the destination storage system.

Likewise, no such disclosure is found in Yanai or Tremblay either. Note that in Tremblay, the “references” which the Examiner cites are references *to particular write requests* from the host, *not* references to previous data transfers between the source and destination storage system; see Tremblay at col. 6, lines 49-50.

In rejecting dependent claim 4, the Examiner contends that “Selkirk teaches a method, wherein each said reference comprises an indication of a location at which the corresponding block was located within the data transfer,” citing Selkirk at page 1, paragraph 10 (Final Office Action, p. 7). Applicants respectfully disagree. Applicants find no disclosure, in the cited section or elsewhere in Selkirk, of a reference that comprises an indication of a location at which the corresponding block was located *within a data transfer*. Likewise, no such disclosure is found in Yanai or Tremblay either.

Therefore, no combination of the cited references discloses *all of the limitations* of independent claims 9, 13, 18 and 21. For the above additional reasons, therefore, claims 9, 13, 18 and 21 and all claims which depend on them are patentable over the cited art.

Independent claim 13

Applicants' arguments above also apply to independent claim 13. In addition, note that claim 13 recites:

13. (Original) A method of mirroring data, the method comprising, in a first storage server:

- receiving a plurality of requests to write data from a set of client devices, the requests for causing modification of a plurality of blocks of data;

- creating a log entry for each of the requests;

- transmitting the log entry for each of the requests to a second storage server located at a secondary site, using one or more data transfers, each of the data transfers including one or more of the modified blocks and having a unique transfer ID;

- saving modified data in a first set of non-volatile storage devices coupled to the first storage server based on the requests; and

- initiating synchronization of data in the first set of non-volatile storage devices with data stored in a second set of non-volatile storage devices coupled to the second storage server, wherein **said initiating synchronization includes**

- for each of the plurality of blocks **which has been only partially modified as a result of the requests**, sending the partially modified block to the second storage server, and

- for each of the plurality of blocks **which has been wholly modified as a result of the requests**, sending a transfer ID and a block number associated with the wholly modified block to the second storage server instead of the wholly modified block, the transfer ID identifying a data transfer in which the wholly modified block was sent to the second storage server during said transmitting the log entry, the block number indicating a location of the wholly modified block within said data transfer. (Emphasis added.)

The Examiner contends that Selkirk discloses the limitations emphasized above in bold, at page 11, paragraph 168; page 5, paragraph 85; and page 6, paragraphs 86-87 (Final Office Action, p. 9). Applicants respectfully disagree.

First, the cited disclosure in Selkirk has nothing to do with data synchronization. Second, Applicants find no disclosure or suggestion in Selkirk of distinguishing between *wholly* modified blocks and *partially* modified blocks, or of performing different actions for wholly modified blocks and partially modified blocks as recited in claim 13. Selkirk does not discuss wholly modified blocks versus partially modified blocks or sending such different types of blocks from one storage server to another. Therefore, no combination of the cited references produces *all of the limitations* of claim 13. For this additional reason, therefore, claim 13 and all claims which depend on it are patentable over the cited art.

Dependent Claims

In view of the above remarks, a specific discussion of the dependent claims is considered to be unnecessary. Therefore, Applicants' silence regarding any dependent claim is not to be interpreted as agreement with, or acquiescence to, the rejection of such claim or as waiving any argument regarding that claim.

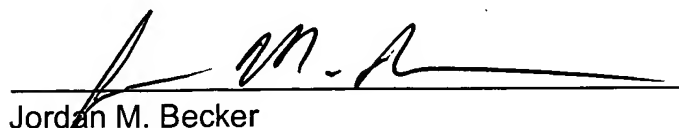
Conclusion

For the foregoing reasons, the present application is believed to be in condition for allowance, and such action is earnestly requested.

If there are any additional charges/credits, please charge/credit our deposit
account no. 02-2666.

Respectfully submitted,
BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP

Dated: October 23, 2006


Jordan M. Becker
Reg. No. 39,602

Customer No. 48102
12400 Wilshire Blvd.
Seventh Floor
Los Angeles, CA 90025
(408) 720-8300